

Comparing Template-Count-Limited Search Strategies for High Mass-Ratio LISA inspirals

James Graber
jgra@mailaps.org

Sixth International LISA Symposium
June 19-23, 2006
Goddard Space Flight Center
Greenbelt, Maryland

Template-Count-Limited Search Equals Selective Search

Not long ago, it was realized that
the search for chirps in LISA's data stream
will be

Computation Limited,
NOT
Noise Limited.

There will be chirps in LISA's data stream with more than adequate SNR
that will not be detected, due to insufficient computing power.

This applies particularly to
Extreme Mass-Ratio Inspirals (EMRIs)
and
Intermediate Mass-Ratio Inspirals (IMRIs)

We will have to select which chirps to search for,
and which chirps to allow to go undetected,
at least initially.

Data Mining LISA's data stream
will be potentially fruitful
many years after completion of the mission.

Many Choices Exist:

Among Them:

Concentrate on a Single Frequency

Concentrate on a Single Time

Concentrate on a Subset of Chirps

The Computational Limit is a Template Count Limit

The signals are buried in noise.

The must be extracted by matched filtering.

Each matched filter search requires a unique template which corresponds to one narrow class of possible signals.

The number of different possible chirp signals is enormous.

Each matched filter search requires computation.

The required computational power and time are not available.

Template Counts and LISA Science Objective No. 5

Lisa Science Objective No. 5
is to perform precision tests of general relativity
by measuring the mass, spin and at least one higher multipole
moment
of a black hole.

This requires an extreme or intermediate mass-ratio inspiral,
EMRI or IMRI.

EMRIs and IMRIs in LISA's sensitive band
will typically have inspirals
with more than 10^5 observed cycles.

Each such inspiral must be measured
to an accuracy of one cycle or less,
both for detection and to perform the science objective.

This requires an accuracy of one part in 10^5 or better.

How many templates are required to adequately represent an EMRI chirp?

One Computation:

Seventeen astrophysical parameters describe a binary inspiral.

All seventeen affect the shape of the EMRI waveform.

Seven or eight parameters contribute significantly.

An accuracy of 10^{-5} is required for each parameter.

10^{35} to 10^{40} templates are possible.

Another computation:

The chirp wave form frequency evolution function (FEF)

can be represented as a Taylor series.

If the first term of the Taylor series

contributes 10^5 cycles to the FEF,

and the second term contributes 10^4 cycles,

and the third term contributes 10^3 cycles,

and the fourth term contributes 10^2 cycles,

and the fifth term contributes 10 cycles,

and the sixth and all higher terms contribute one cycle or less,

more than 10^{15} templates will still be required.

In a realistic LISA EMRI chirp,

all of the above numbers will be exceeded.

How many templates can be searched?

Gair et al have estimated as many as 10^{12} by the time LISA flies,
based on Moore's law increases.

Others have estimated that 10^{10} would be possible now.

Matching Partial Chirps

We cannot match the entire possible chirp set with a computationally feasible set of parameters

We can only match as much of a chirp as can be described by a smaller number of parameters which combine into a computationally feasible set.

This means we must describe partial chirps.

This means our integration time will be limited.

Gair et al determine that the likely limit is three weeks or less, rather than the previously considered one year integration time.

This means 10^3 to 10^4 cycles, instead of 10^5 .

In the context of the Taylor series model,
it means fewer terms contribute,
as well as fewer cycles for each term.

Effects of Short Integration Times

Loss of Signal-to-Noise Ratio (SNR)

Need to Analyze Partial Chirps

Simpler Chirp fragments

Potential Redundancy

(If you analyze every frequency in every time slice,
you will rediscover each chirp many times.)

Reducing Template Counts in the Taylor Series Model

When an EMRI chirp is represented as a Taylor series
(in $1/r$ or an equivalent variable)
obtained by perturbative or post-Newtonian methods,
it is possible to reduce the template count
for describing a fraction of the chirp in two ways:
by looking at a shorter section of the chirp,
and by looking at an earlier section of the chirp.

Looking at an earlier section of the chirp,
where $1/r$ is smaller,
decreases the number of cycles contributed by the higher order
terms,
and hence decreases the template count.

This decrease will sometimes allow
a larger number of cycles to be contributed
by the leading order term,
and still have a computationally feasible template count.
This increase can allow a longer integration time,
and hence a higher SNR.

In many cases the net result is that
the chirp is easier to detect
at the earlier stage of the inspiral,
because of the lower template count,
and the higher SNR.

Other Advantages of Early Chirp Detection and Analysis

In addition to
Simpler Signals
with fewer parameters
and (sometimes) longer integration times
and (sometimes) higher SNR

Early chirp phases where $1/r < .025$
have more robust theoretical predictions
which may be necessary
to complete LISA Science Objective No. 5,
the precision testing of general relativity.
(See gr-qc/0604080)

A New Strategy For EMRI and IMRI Searches With Feasible Template Count.

Strategic Approach:

1. Spend most computational resources on an initial search. (i.e. 90%)
Search for as many candidate chirps as computationally possible.
Accept approximately one tenth of one percent, i.e one of a thousand.
(Three sigma)
Note that most accepted candidates at this stage will be false alarms.
Concentrate resources on a single frequency or at most two frequencies.
(Entrance and possibly exit to LISA's most sensitive band, .003 hz and .01 hz)
2. Continue Search in a coherent manner (Find and Follow, See below)
(Contrast Gair et. al. who search all frequencies and assemble incoherently)
Extend Search for roughly equal time and template count.
Consider less than 100 follow up options for each primary chirp candidate.
At each step either refine previous parameters or introduce one new parameter.
Accept less than ten percent of candidates from previous stage.
3. Iterate Step 2 until false alarm probability is $\ll 1$, ie about ten times.
The remaining candidates are genuine chirps with high probability.
4. By now, parameters have been increased and refined.
A combined integration is possible, as well as a precision analysis.

Finding Initial Chirp Candidates

Select EMRI and IMRI chirps to search for.

Subdivide into groups by chirp mass and mass ratio.
Apportion template count allotments to groups.

For each group,
determine if low frequency search or high frequency search
is more economical.

Prepare templates for each group
either low or high frequency,
whichever is most economical,
up to maximum length allowed by template count allocation..

Perform standard matched filter searches.

(Typically, all EMRIs will generate low frequency searches.
Some IMRIs may generate high frequency searches.)

If 10^{12} total templates are possible,
each group could have of order 10^{10} templates.

This might allow 10,000 cycle integrations for some IMRIs,
but could require some EMRIs to be integrated for less than 500
cycles

Initial Search: Taylor Series Analysis

Assuming a template count allocation of about 10^{10}
for a group characterized by chirp mass and mass ratio,
or equivalently large and small masses,
a classical EMRI in the LISA band,
i.e. neutron star spiralling into a 10^6 solar-mass supermassive black hole,
might have an initial chirp fragment 300 cycles long
with the second term contributing 100 cycles,
the third term 50 cycles, the fourth term 20 cycles,
the fifth term 10 cycles, the sixth term 5 cycles,
the seventh term 2 cycles and
the eighth and higher terms contributing
less than one cycle each.

(This example has been deliberately over simplified and over rounded.)

Similarly, an IMRI with a 10^2 solar-mass black hole
spiralling into a $10^{4.5}$ solar-mass intermediate-mass black hole
might have an initial chirp fragment 5000 cycles long
with the second term contributing 1000 cycles,
the third term 100 cycles, the fourth term 10 cycles,
and the fifth and higher terms contributing
less than one cycle each.

Actually, it is the variance in the number of cycles contributed,
rather than the number of cycles contributed,
which enters as a factor in the template count.
However, these two numbers are usually of similar magnitude.

Note that each of the above two examples requires only about 10^6 templates
in the single frequency search mode,
but the first example would require about 10^8 templates
and the second example would require about 10^{10} templates
if we were searching all frequencies.

Net Result: We will know three to seven
Taylor Series Coefficients

Coherent Search Continuation: Find and Follow

Once we know the initial coefficients of the Taylor series, we can refine them and introduce one more at the next lower order by introducing only two or three alternatives for each coefficient.

Thus the number of searches (new templates) at each refinement step will be of order 100 or less.

Since only one of a thousand candidates was accepted in the first step, the number of searches in the second step will be less than 10% of the number of searches in the first step.

Since only 10% or less of candidates are accepted in the second and subsequent steps, the number of searches in the third and subsequent steps will be less than 10% of the number of searches in the preceding step.

This reduction will continue until all false candidates are rejected, in which case the search case terminates with zero candidates, or continues with only one candidate which is a genuine chirp.

Comparison of Results:

Detection rate

Previously, it was hoped that chirps could be integrated for one year.

Currently, three weeks is thought to be a better estimate.

Previously it was thought that an eight- to ten-sigma result would be needed for one years integration to guarantee an adequately low false alarm rate.

Using the standard square root law for SNR, a sigma of eight for one year integration is equivalent to a sigma of four for thirteen weeks integration and a sigma of two for about three weeks integration.

The proposed method requires a sigma of three for the first three week integration, and will have a sigma exceeding eight if a full years follow on tests succeed. So the false alarm rate of the new method exceeds that of the previously hoped for technique, now known to be impractical for template count computational reasons.

But the new method will find fewer chirps because of the requirement of three sigma, rather than the equivalent of two for the first three weeks. The difference between two and three sigma is about equivalent to a distance ratio of two to three or a volume ratio of 8 to 27. This is a significant, but not devastating loss.

Comparison of Results: Template Counts

Since the LISA sensitive band runs from .003 hz to .01 hz, the template count for a full frequency search over this entire band is roughly 70% of that for a single frequency search at the high end multiplied by the number of cycles in that search.

Thus for a search 1000 cycles long, we save a factor of 700 on a single frequency search, and a factor of 350 on a double frequency search.

This is a significant savings.

References and Acknowledgements.

This type of search is based on the single dominant frequency of the chirp.

It is my recollection that the single frequency approach was mentioned to me by Sam Finn at a coffee break at Pen state about five years ago.

I thank Sam for pointing out this possibility to me.
I have found no published references to this approach.

References

Finn and Thorne gr-qc/0007074

Gair et. al. gr-qc/0405137

Graber gr-qc/0604080